arXiv:1112.5121v1 [physics.soc-ph] 21 Dec 2011

COLLABORATION NETWORK FORMATION AND THE DEMAND FOR PROBLEM SOLVERS WITH HETEROGENOUS SKILLS

KATHARINE A. ANDERSON
CARNEGIE MELLON UNIVERSITY, TEPPER SCHOOL OF BUSINESS

Abstract. Collaboration networks provide a method for examining the structure of collaborative communities. The model I present in this paper connects an individual's skill set to her position in the collaboration network, and changes in the distribution of skills to the structure of the collaboration network as a whole. I show that individuals with a useful combination of skills will have a disproportionate number of links in the network, resulting in a skewed degree distribution. The degree distribution becomes more skewed as problems become more difficult, leading to a community dominated by a few high-degree superstars.

Heterogeneity is important in a wide variety of creative and knowledge-based fields, including product development, art and design, entrepreneurship, and academic research. Innovation in these areas is driven by diverse groups of people, working together on difficult problems. In that context, simple models with only a one-dimensional ability level cannot be used to study the institutional arrangements that are driven by and respond to the multidimensional factors that make people different. This kind of complex heterogeneity has garnered increasing interest in recent years, in part because it is not well understood. In this paper, I present a framework for understanding heterogeneity in a context where it is not peripheral, but central: collaborative production.

Thanks to Scott Page and Ross O'Connell. This work was completed with the help of funding from the NSF and computing resources from the Center for the Study of Complex Systems.

This allows me to examine the link between rich individual heterogeneity, and the complex structure of collaborative communities.

Heterogeneity is particularly central to collaborative production, because while some common ground is required to facilitate communication, complementarities provide the real advantage. Collaboration works when individuals with diverse skills come together to solve problems that none of them could solve alone.[1] Collaboration networks provide a way of looking at the structure of collaborative communities. The nodes in the network represent individuals (eg: academics, researchers, policy-makers, or entrepreneurs) and two individuals in the network are linked if they have worked together on a problem. In this paper, I present the first model of the formation of collaboration networks that incorporates individual heterogeneity. I start with a population of problem solvers, each of whom is endowed with a set of skills. The individuals collaborate on a set of problems, and the result is a collaboration network. By drawing direct connection between the distribution of skills in the population and the structure of the collaboration network, I am able to look at how the nature of a collaborative community (as measured by the distribution of skills in that community) affects the structure of the collaboration network as a whole.

The structure of this network is important for a number of reasons. On an individual level, a person's position on the network reflects her role in the collaborative community–a person who is located on the periphery of a community will play a different role than a person who is more central. A person's position in the network is also correlated with her influence in the collaborative community,[2] her access to knowledge, tools, and information,[3] and her probability of finding a job.[4] On a more global level, the overall structure of the collaboration network reflects the nature of the problem solving community–a community consisting of several densely-connected clusters will have a different character than a community where links are spread evenly. Moreover, network structure impacts the function of the community, improving or impeding the flow of information and ideas, opinion formation, and the spread of effective technologies.[5]

---

[1] Hong and Page (2001), Polzer et al. (2002), Thomas-Hunt et al. (2003), Phillips et al. (2004)

This suggests two interesting classes of questions relating to collaboration networks:

(1) What characteristics of an individual affect her position on the collaboration network?
(2) What aspects of the collaborative process (eg: institutions, characteristics of problems and problem solving communities) affect the overall structure of the collaboration network?

In order to address these questions, I need a model of how collaboration networks form, and, in particular, a model of how individual heterogeneity translates into network heterogeneity.

Existing network formation literature can be divided roughly into two categories: statistical models and decision-based models. In statistical models, the process of link formation is stochastic. Examples of these types of models include Erdos-Renyi random networks[6], Watts-Strogatz small world networks,[7] incumbency models[8], and preferential attachment models and their derivatives.[9] These models of network formation have the advantage of being mathematically easy to work with. Preferential attachment and incumbency models have the additional advantage that they can be made to reproduce the observed distribution of degree, with it's extremely long right tail. However, because links in these networks are made randomly, they do not allow us to answer questions about incentives and network structure. Moreover, they do not incorporate any true individual heterogeneity, and thus cannot be used to suss out which individual characteristics impact an individual's position in the network.[10]

---

[2] see Menzel and Katz (1955), DeMarzo et al. (2003), Golub and Jackson (2010), Banerjee et al. (2011)

[3] see Jackson (2008) for a summary of this literature

[4] See Calvó-Armengol (2004) and Calvó-Armengol and Jackson (2004)

[5] see Jackson and Yariv (2007), Jackson and Rogers (2007) for examples, plus a survey of the literature

In decision-based models, individuals in the network choose their network connections in a way that maximizes an objective function. For example, in the Connections Model,[11] players gain a benefit for both direct and indirect links (the benefit to indirect links decays across multiple links) and pay a cost for each direct link made. In the Coauthor Model,[12] players are must allocate their effort across multiple projects where the payoff from a paper is inversely related to the number of links the two coauthors have. In Goyal and Moraga-González

---

[6] Erdös and Rényi (1959)

[7] Watts and Strogatz (1998)

[8] Guimerà et al. (2005)

[9] See, among others, Barabási and Albert (1999), Jackson and Rogers (2007) and Ramasco et al. (2007)

[10] For example, in a preferential attachment model, new nodes connect to existing nodes with a probability proportional to the degree of the existing node. This yields a long-tailed degree distribution, where a small number of individuals have the majority of the links, while most individuals have very few. However, their are only two factors that influence which nodes will have the highest degree in the network: age (older nodes gain an early advantage, which compounds over time) and luck. Thus, if you wanted to look at a community and predict which nodes will have high degree, this type of model would not be useful

[11] Jackson and Wolinsky (1996)

[12] Jackson and Wolinsky (1996)

(2001), firms choose a set of links and an effort level to put into research and development, given that such links general both perfect and imperfect spillovers. These types of models provide more insight into the link between incentives and network structure than purely statistical models. However, in all of these models, the individual agents are homogeneous, and thus the resulting networks are symmetric. Because individuals in these models are homogeneous, they can't tell us anything about the heterogeneity of network position we observe empirically. Moreover, the resulting networks don't look much like empirical social networks, and thus they can tell us only a limited amount about global network structure.

The model I present in this paper is decision-based, and it is the first such model of network formation to incorporate individual heterogeneity. Each individual in the model has a set of skills, drawn from a larger pool.[13] That set of skills represents the human capital of the problem solver–the tools, techniques and knowledge that are useful inputs to knowledge-based production. Any individual may have a subset of those skills, but it is unlikely that a single person will have all of them. Each individual in the community is endowed with their own problem and the collaboration network forms when individuals find collaborators to help them solve those problems.

Using this model, I show that an individual's position on the network is a highly non-linear function of her set of skills. In particular, the number of connections an individual has on the network (her degree) is a supermodular function of her set of skills. In other words, an individual with a useful combination of skills will have a disproportionate number of links on the collaboration network. Thus, the relationship between an individual's set of skills, and her importance to the collaborative community cannot be captured by a linear pricing structure.

---

[13] This idea of giving individuals multiple skills is rooted in an older literature, starting with Roy (1951), which gives individuals an ability level in each of two different occupations (for example, Roy (1951) gives an individual a skill level as a "hunter" and a "fisherman". This, in essence, is equivalent to assuming that individuals have ability levels in two different skills. There has been recent work that explicitly starts to migrate towards a model where individuals can utilize both skills for tasks (rather than simply using one or the other). Examples include Lazear (2004, 2005) and Astebro and Thompson (2011). However, all of these works confine their analysis two skills. Here, I consider skill sets of arbitrary size.

The importance of combinations of skills has a dramatic affect on the overall structure of the collaboration network. Using a special case called the Bernoulli Skills Model, I show that even when skills are distributed independently, the distribution of degree in the network will be highly skewed–a few individuals will have a large number of links, while the vast majority of individuals will have very few. Small initial differences in individual skill sets create large differences in the distribution of links, meaning that the distribution of links in the collaboration network may be skewed, even when the distribution of skills is not. This skewed degree distribution closely matches the distribution that we observe in empirical collaboration networks.

Moreover, because the model creates a direct link between the skill population and the collaboration network, I can use comparative statics to make predictions about how changes in the population of problem solvers will change the structure of the network as a whole. Using the Bernoulli Skills Model, I show that as problems become more difficult, the degree distribution of the network becomes more skewed, and the network becomes dominated by a few, high-degree superstars. Using a second case called the Ladder Model I observe another pattern–the network becomes increasingly skewed as skills become increasingly hierarchical.

Finally, it is worth nothing that the framework and results I present here have broader implications for our understanding of heterogeneity in labor markets. I conclude with a brief discussion of these issues.

1. A General Model of Skills, Problem Solving, and Collaboration Networks

1.1. **Inputs: Problem Solving Population and Problems.** Let $I = \{1, 2, ...N\}$ be the set of problem solvers (individuals).

Let $S = \{a_1...a_M\}$ denote the (finite) set of all skills.

A individual i's *skill set* is the subset of those skills she possesses, $A_i \subseteq S$. The frequency of skill set A in the population is given by $\Psi(A)$, a probability measure with support $\Sigma(\Psi) \subseteq 2^S$–that is, $\Psi(A)$ is the fraction of the individuals in I who have the skill set $A \subseteq S$.[14]

Each individual is endowed with a copy of a problem requiring a subset of the skills available in the population, $\omega \subseteq S$.[15]

Thus, the inputs to the model are a set of skills used to solve problems (S) and a population of problem solvers (Ψ).

1.2. **Collaboration and Problem Solving.** A *collaboration* is a subset of the problem solvers, $C \subset I$. A collaboration can *solve* a problem if together they possess all of the required skills–that is, if the problem is solved if $\omega \subseteq \bigcup_{j \in C_i} A_j$.

The problem yields a payoff of 1 if solved. If a problem solver can solve her problem alone (that is, if $A_i = \omega$) then she keeps the entire payoff. If she solves it with the help of others, she splits the payoff evenly with her collaborators, giving each a share of $\frac{1}{|C_i|}$.[16]Since others in the community face their own problems, a problem solver may be asked to help with other problems as well. Thus, individual i's payoff is the sum the payoff she gets from solving her own problem, plus the payoffs she gets from collaborating with others on their problems:

---

[14] Formally, Ψ is a *frequency* distribution, rather than a probability distribution–that is, Ψ is a *realized* distribution of skill sets across the population. The distinction between frequency and probability distributions disappears when N is large–using a frequency distribution allows me to also make statements about small N as well.

[15] For simplicity, I assume that all individuals face the same problem. The results are identical if each individual faces a different problem drawn from a known distribution of problems, Ω, with support 2^S, and collaborators are chosen *ex ante*. If they choose their collaborators *ex post*, then the results are similar, if notationally more complex.

[16] This particular distribution scheme has a number of points to recommend it: it does not require that agents have control over the way payoffs are split (advantageous when payoffs

$$u_i = \frac{1}{|C_i|} + \sum_{j \neq i \, st \, i \in C_j} \frac{1}{|C_j|}$$

[17] An individual chooses her set of collaborators (C_i) to maximize her utility. Note that her payoff to solving her own problem is always positive, and thus it is always incentive compatible for her to solve the problem. Since each individual controls only her own collaborative decisions, the utility-maximizer chooses C_i to minimize the number of collaborators she must work with on her own problem–in other words, she chooses a minimal subcover of the set of skills she lacks–$A_i^c = \omega_i \backslash A_i$. Let $\mathbb{C}_i$ denote the set of all minimal subcovers of A_i^c. If there exist multiple minimal subcovers (ie: if $|\mathbb{C}_i| > 1$) then I assume that the individual chooses $C_i^* \in \mathbb{C}_i$ at random.[18]

It is worth noting that while I have chosen to model the motivate the decision to minimize the set of collaborators via an equal split of payoffs, the results that follow are unchanged under any distribution scheme which induces problem solvers to minimize the number of collaborators they use. So, for example, if collaborating with additional individuals is costly (eg: in terms of time or communication), then the results that follow will still hold with no

---

are non-monetary), and it gives each individual her Shapely value for the coalition. Many alternative splitting schemes will be behaviorally identical (see below).

[17] It is important to highlight that this payoff is not equivalent to a market wage. A model producing equilibrium wages would require considerable additional machinery, including time constraints and search costs, and is therefore beyond the scope of the current work.

[18] Since the individual indifferent between minimal subcovers, this choice at random follows convention. The results that follow are not sensitive to this assumption. In particular, if an individual is biased towards choosing collaborators who are not as busy (eg: those with fewer links) then the results that follow hold precisely, rather than on average. If an individual is biased towards choosing collaborators who are popular (eg: have many links) than the results that follow are qualitatively similar, and if anything are more striking.

modification. The results will not hold under any payoff structure where it is not individually rational to minimize the set of collaborators used to solve a problem.

1.3. Complementary Skills Networks.

For a given a set of collaborations, $C = \{C_1 ... C_N\}$, the collaboration network is represented by an adjacency matrix, $g\left(C\right)$, where $g_{ij}\left(C\right) = 1$ if $j \in C_i$. Note that the network is directed–since $j \in C_i$ does not necessarily imply $i \in C_j$, it may be that $g_{ij}\left(C\right) \neq g_{ji}\left(C\right)$. However, the links are mutual, in the sense that j will never want to terminate a link (see Section 1.6.1 for further discussion). When all collaborators are chosen optimally (that is, when $C_i \in \mathbb{C}_i \, \forall i$), I will call the result a *complementary skills network* (or *complementarity network* for short).

Definition. A network, $g\left(C\right)$, is a *complementary skills network (complementarity network)* if each individual in the network chooses a minimal set of collaborators required to solve her problem–eg: if $C_i \in \mathbb{C}_i \, \forall i$.

Since the set of minimal subcovers each individual ($\mathbb{C}_i$), depends on the distribution of skills in the population, I use $\Gamma\left(\Psi\right)$ to denote the set of complementary skills networks for a particular distribution of skills, Ψ. In other words, $\Gamma\left(\Psi\right)$ is a set of networks over which all individuals are maximizing utility. In all of the measures that follow, I will average over all networks in $\Gamma\left(\Psi\right)$.[19]

1.4. Network Measures.

Recall that I will be looking at two different questions: first, how an individual's position on the network depends on her set of skills, and second how the overall structure of the network depends on the distribution of skills in the population. For both of these questions, I will focus on degree: the number of connections an individual has on the network.

---

[19] Ironically, the great strength of social networks (their rich heterogeneity) also makes them difficult to work with. One method for simplifying this problem is to divide networks into equivalence classes of networks that are similar in some meaningful way. Here, the networks are divided into equivalence classes according to the individual agents' indifference condition. That this division is not arbitrary–but rather a nature result of the decision rule–is one of the great strengths of this particular model.

First, I will look at how an individual's average in-degree on the networks in $\Gamma(\Psi)$ depends on her skill set. I will denote individual i's in-degree on a particular network as d_i. Her average in-degree will be $E[d_i]$, where the expectation is taken over all networks in $\Gamma(\Psi)$–that is, all networks that the individual is indifferent between. In Section 2, I will consider the mapping between an individual's skill set an her expected in-degree in a particular problem-solving population, $E[d_i] = f(A_i, \Psi)$.

I will also consider how the overall distribution of in-degree depends on the distribution of skills in the population. Let Δ denote the distribution of expected in-degree. That is, $\Delta(d)$ is the fraction of the individuals who have expected in-degree d, where the expectation is taken over all $g(C) \in \Gamma(\Psi)$.[20]

Before continuing, a brief word about network notation is in order. First, note that for ease of reading I will usually drop the argument of $g(C)$. I will denote a link from i to j by ij. Using a slight abuse of notation, I will use g to refer to both the adjacency matrix (as above) *and* the set of links in the network–that is, $ij \in g$ if i is connected to j in the network g. In a similar abuse of notation, I will use $g - ij$ to represent the network that results when the link ij is removed from an existing network, g, and $g + ij$ to represent the network that results when the link ij is added to the existing network, g. Finally, for clarity in the exposition, I will refer to in-degree simply as "degree".[21]

---

[20] Alternatively, we might plot the distribution of degree across all networks $g \in \Gamma(\Psi)$. That is, we could set $\Delta(d) = \sum_{g \in \Gamma(\Psi)} \delta_g(d)$ where $\delta_g(d)$ is the fraction of nodes in network g with degree d. This choice does not affect the results.

[21] This should cause no significant confusion. I use in-degree because it has a clear, empirical interpretation, but the results qualitatively similar if we consider an individual's degree in

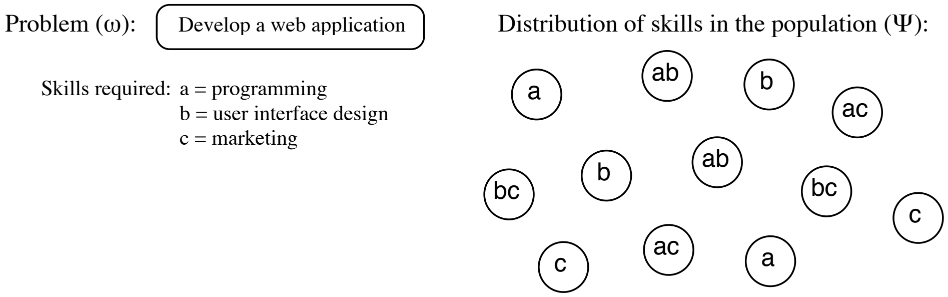
Problem (ω):
Develop a web application
Skills required: a = programming
b = user interface design
c = marketing
Distribution of skills in the population (Ψ):
a
ab
b
ac
bc
b
ab
bc
c
c
ac
a

Figure 1.1. An example of a problem and population of problem solvers

1.5. **Example.** An example will help clarify the structure of this model. Suppose all of the individuals in the population face the same problem requiring three skills: $\omega = \{a, b, c\}\ \forall i$ –for example, developing a web application might require programming skills, user interface design skills, and marketing skills. Suppose the distribution of skills is such that everyone in the population has at least one skill, but no one has all of the skills required. In other words, every individual in the population has something to add, but none of them can solve the problem on their own. Suppose further that each skill combination is equally likely–in other words, $\Psi(A) = \frac{1}{6}$ for all $A \in \{\{a\}, \{b\}, \{c\}, \{ab\}, \{ac\}, \{bc\}\}$ and $\Psi(A) = 0$ otherwise (see Figure 1.1). In this particular population of problem solvers, each individual needs exactly one collaborator to solve the problem–that is, someone with skill set $\{a\}$ must link to someone with the skill set $\{b, c\}$, someone with skill set $\{b, c\}$ may choose from those with skill sets $\{a\}, \{a, b\}$, and $\{a, c\}$, and so on. The problem solver is indifferent between any two individuals who have the skills that she needs. By linking those who collaborate on problems, we obtain a collaboration network. There will be one such network for every set of optimal collaborations. Figure 1.2 shows two collaboration networks for this population of problem solvers and two optimal

---

the directed network (the sum of his in-degree and out-degree), or use the degree of the individual in a network where directed links are projected into undirected links.

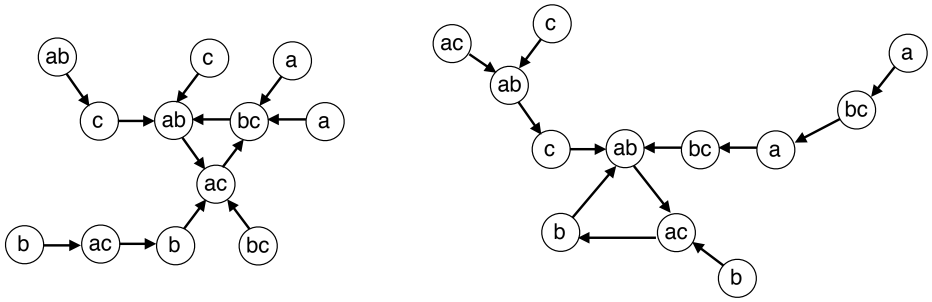

FIGURE 1.2. An example of two networks formed from population/problem in Figure 1.1

choices of collaborators. The set of all such networks for a given population of problem solvers is denoted by $\Gamma(\Psi)$.

1.6. **Discussion.** This model produces outcomes consistent with several empirical facts about collaboration and problem solving. First, the model predicts that collaboration will be more important when the average individual has a thinner slice of the total skills required to solve a problem–in other words, as problems become more difficult, collaboration networks will become more densely connected. This prediction is born out in the data–collaborative work has become increasingly common in a variety of academic fields.[22] Moreover, the literature supports a connection between this increased reliance on collaboration and the difficulty of problems faced, specifically the increasing complexity of required methodologies.[23] This highlights one of the advantages of using this model of collaboration network formation: statistical models of network formation–including preferential attachment–assume constant degree over time, and thus cannot be used to make predictions about the relationship between problem difficulty and collaborative effort.

---

[22] See Barabasi et al. (2002) (physics), Grossman and Ion (1995) (mathematics), Moody (2004) (sociology), Acedo et al. (2006) (management science), Laband and Tollison (2000) and Goyal et al. (2006) (economics).

[23] Laband and Tollison (2000)and Moody (2004)

1.6.1. *A Note on Stability and Efficiency of the Complementarity Network.* Before considering any specific questions about the complementarity network, it is worth considering it's stability and efficiency. Jackson and Wolinsky (1996) introduce an equilibrium concept of network stability, called *pairwise stability.* Briefly, a network is pairwise stable if no individual would prefer to terminate an existing link, and if no pair of individuals would prefer to add a link (see Appendix A.1 for a more formal definition in the case of a directed network). Pairwise stability implies that links are mutual, because both individuals involved agree to maintain the link. Theorem 1 states that any complementary skills network is pairwise stable, implying that all links in the network are mutual.[24] Moreover, any complementary skills network is strongly efficient–that is, the population extracts the maximum possible value from the network.[25]

Theorem 1. *Any complementary skills network, $g \in \Gamma(\Psi)$, is pairwise stable and strongly efficient.*

Proof. See Appendix A.1 □

2. Skills and Degree in the Collaboration Network

In the context of a collaboration network, an individual's degree represents the number of collaborators she has, and is a proxy for her importance in the collaborative community. In this section, I consider the relationship between an individual's skill set, and her degree in the collaboration network.

---

[24] Because the equilibrium links are mutually beneficial, one could functionally think of a complementarity network as either directed (because that is how it is constructed) or undirected (because that it is how it functions). For ease of reading, I will usually omit the directional arrows from networks pictured in this paper.

[25] This result–that equilibrium networks are efficient–contrasts with two other models of social network formation–Jackson and Wolinsky (1996) and Goyal and Moraga-González (2001)–in which pairwise stable networks tend to have more links than is efficient.

2.1. **Skills and Degree: An Example.** Before presenting general results, it is useful to see an example. Consider the example from the previous section: each individual faces a problem requiring three skills, $S = \{a, b, c\}$, and they each have one or two of those skills, so that Ψ has support $\{\{a\}, \{b\}, \{c\}, \{ab\}, \{ac\}, \{bc\}\}$.

An individual's in-degree on the network will depend on the number of people who need her skills (the demand for a subset of her skills) and the number of other people who can provide those skills (the supply of that subset of skills). For example, consider someone with the skill set $\{a\}$. Her skills are in demand by anyone who needs skill a–in this example, anyone with skill set $\{b, c\}$. A person with skill set $\{b, c\}$ may collaborate with anyone who has skill a, including those with skill sets $\{a\}, \{a, b\}$, or $\{a, c\}$. So in this example, the expected degree of an individual with skill set $\{a\}$ is

$$E\left[d\left(\{a\}\right)\right] = \frac{\Psi\left(\{b,c\}\right)}{\Psi\left(\{a\}\right) + \Psi\left(\{a,b\}\right) + \Psi\left(\{a,c\}\right)}$$

Similarly, an individual with the skill set $\{a, b\}$ can help anyone who needs skill a, anyone who needs skill b, and anyone who needs skills a and b, yielding expected degree

$$E\left[d\left(\{a,b\}\right)\right] \quad = \quad \frac{\Psi\left(\{b,c\}\right)}{\Psi\left(\{a\}\right) + \Psi\left(\{a,b\}\right) + \Psi\left(\{a,c\}\right)} + \frac{\Psi\left(\{a,c\}\right)}{\Psi\left(\{b\}\right) + \Psi\left(\{a,b\}\right) + \Psi\left(\{b,c\}\right)} + \frac{\Psi\left(\{c\}\right)}{\Psi\left(\{b,c\}\right)}$$

Note that an individual with skills a and b will have more links, on average, than an individual with skill a and an individual with skill b put together: $E\left[d\left(\{a,b\}\right)\right] > E\left[d\left(\{a\}\right)\right] + E\left[d\left(\{b\}\right)\right]$. This is because an individual with both skills can help anyone who needs skill a, anyone who needs skill b, *and anyone who needs both.* This means that an individual with both of those skills will have, on average, more than twice the number of links than a person with either of those skills in isolation.

2.2. **Skills and Degree: General Results.** A similar type of result holds more generally. Theorem 2 states that an individual's expected degree in a complementary skills network is a supermodular function of her set of skills. That means that regardless of the skill set required for the problem or the distribution of skills in the population, an individual with skill set $A \cup B$ will have at least as many links as individuals with skill sets A and B put together. The intuition for this result is similar to the above example–the set of all problems that can be solved by someone with the skill set $A \cup B$ includes those that can be solved by someone with skill set A, and those that can be solved by someone with skill set B, plus those requiring some skills from *both* sets.

Theorem 2. *For any set of skills, S, and distribution of those skills, Ψ, an individual's expected degree over the networks in $\Gamma(\Psi)$ is a supermodular function of her set of skills. That is, $Ed(A \cup B) + Ed(A \cap B) \geq Ed(A) + Ed(B)$.*

Proof. Here, I will prove the result for the case where an individual needs only one collaborator to solve her problem. For the sake of clarity, also I consider the case where $A \cap B = \emptyset$. The proof for the general result is similar, and can be found in Appendix A.2. Since $d(A \cap B) = d(\emptyset) = 0$, we need to show that $Ed(A \cup B) \geq Ed(A) + Ed(B)$. Consider $d(A \cup B)$. The fraction of the population needing $C \subseteq A \cup B$ is $\delta(C) = \Psi(S \backslash C)$. The fraction who can supply the set C is $\sigma(C) = \sum_{D \subseteq S \backslash C} \Psi(C \cup D)$. Thus

$$E[d(A \cup B)] = \sum_{C \subseteq A \cup B} \frac{\Psi(S \backslash C)}{\sum_{D \subseteq S \backslash C} \Psi(C \cup D)} = \sum_{C \subseteq A \cup B} \frac{\delta(C)}{\sigma(C)}$$

We can partition $C \subseteq A \cup B$ into three categories:

1) $C \subseteq A$

2) $C \subseteq B$

3) $C \subseteq A \cup B$, $C \nsubseteq A$, $C \nsubseteq B$.

This division gives us the following:

$$\begin{aligned} E\left[d\left(A \cup B\right)\right] &= \sum_{C \subseteq A} \frac{\delta\left(C\right)}{\sigma\left(C\right)} + \sum_{C \subseteq B} \frac{\delta\left(C\right)}{\sigma\left(C\right)} + \sum_{C \subseteq A \cup B \, and \, C \cap A,\, C \cap B \neq \emptyset} \frac{\delta\left(C\right)}{\sigma\left(C\right)} \\ &= E\left[d\left(A\right)\right] + E\left[d\left(B\right)\right] + \phi \\ &\geq E\left[d\left(A\right)\right] + E\left[d\left(B\right)\right] \end{aligned}$$

□

This theorem suggests an immediate corollary.

Corollary 3. *Adding skills to an individual's skill set will never decrease her degree in a cost minimizing collaboration network.*

Proof. From Theorem 2, $E\left[d\left(A \cup a\right)\right] + E\left[d\left(A \cap a\right)\right] = E\left[d\left(A \cup a\right)\right] \geq E\left[d\left(A\right)\right] + E\left[d\left(a\right)\right]$, and so $E\left[d\left(A \cup a\right)\right] - E\left[d\left(A\right)\right] \geq E\left[d\left(a\right)\right] \geq 0$. □

2.3. **Skills and Degree: Discussion.** The result presented in Theorem 2 has some implications for who we think is important in a collaborative community. In particular, it indicates that it is important to consider an individual's *combination* of skills, rather than looking at each of her skills in isolation. For example consider the population summarized in the first panel of Figure 2.1. The problem requires five skills, $S = \{a, b, c, d, e\}$, which are distributed across a population of $N = 5n$ as indicated. Each skill is held by exactly $2n$ individuals, and thus no skill is rarer than the others. Traditionally, we might condense the information contained in this table into a single-dimensional measure of ability–individuals 1 and 2 have the most skills, and therefore, we might expect them to have the most links. However, despite having fewer skills, individual 3 will, on average have more links. This is because while neither of individual 3's skills are rare by themselves, in combination, they are both rare, and

Population 1

	a	b	c	d	e	In-degree
Individual 1	X	X	X			1.5
Individual 2	X	X	X			1.5
Individual 3				X	X	3
Individual 4				X		0.5
Individual 5					X	0.5

Population 2

	a	b	c	d	e	In-degree
Individual 1	X	X	X			1.5
Individual 2	X	X	X			1.5
Individual 3				X	X	3
Individual 4	X	X		X		0.5
Individual 5	X	X			X	0.5

FIGURE 2.1. Two populations of problem solvers with five skills $S = \{a, b, c, d, e\}$–n copies of each individual yields a population of size $N = 5n$.

useful to a large fraction of the population. This means that agent 3 will have more links than a tally of her individual skills might predict—-in other words, the value of a combination of skills may be greater than the sum of its parts.[26] This further suggests that in problem solving communities such as those in knowledge-based industries, entrepreneurial firms, and academic research, models in which individuals are scored on a one-dimensional ability scale will fail to capture the full effect of variation between individuals.

Pushing this point a bit further suggests another implication of Theorem 2– because degree is a supermodular function of a problem solver's complete set of skills, it is not generically possible to assign prices to individual skills in a way that captures her degree on the social network. This means that examining the supply and demand of single skills in isolation does not necessarily capture an individual's value to a community of problem solvers.

Corollary 4. *There need exist no vector of prices,* μ, *such that* $\sum_{a \in A} \mu_a = d(A)$.

---

[26] Examples of this phenomenon are not difficult to find. For example, consider a tech entrepreneur who has experience in both computer programming and marketing. Neither of those skills are rare individually, but together they are quire rare.

To further emphasize this point, consider the skill distribution shown in the second panel of in Figure 2.1. This distribution is identical to that in the first panel, except that individuals 4 and 5 have been given an extra skill. However, gaining these skills does not influence the degree of either. In fact, endowing them with those skills does not change any part of the degree distribution. Skills a and b have value to individuals 1 and 2, but not to individuals 3 or 4–clearly, no linear weighting of the individual skills could produce such a pattern. Two characteristics of this model contribute to this result. First, skills are bundled within a person. Thus, in evaluating a collaborator, that person's combination of skills must be considered as a unit. Second, the individuals in the model have an incentive to minimize the number of collaborators they work with. Together, these two factors mean that an individual's value to the collaborative community may be more than the sum of her individual skills.

These results have implications for empirical models. If we assume that skills contribute to outcomes independently, then we will underestimate the value of particular combinations of skills and overestimate the value of other combinations. Moreover, this non-linear relationship between skills and degree means that if we observe skills imperfectly, the amount of variation in outcomes that is explained by an observed set of skills will decline dramatically in the set of skills we can observe. Additionally, the fact that an individual's degree depends on both the distribution of skills in the population and the set of skills she already has suggests that in collaborative fields, optimal training decisions are highly individualized.[27]

3. The Distribution of Skills and the Structure of the Collaboration Network: The Bernoulli Skills Model

The degree distribution of a collaboration network affects the function of the collaborative community. A skewed degree distribution indicates that a small number of individuals are involved in most of the collaborations, while most people are involved in very few. In this section, I look at how the distribution of skills in the population affects the degree distribution of the collaboration network.

3.1. The Bernoulli Skills Model.

Thus far, I have considered results for a general skill population. However, when we turn to the effects of skill populations on network structure, it is difficult to obtain clear predictions using such a general construction. Therefore, in this section I will consider a special case, where skills are distributed independently with equal probability–that is, $Prob\left(a_i \in A | a_j \in A\right) = Prob\left(a_i \in A\right) = p \, \forall \, i \neq j \in S$. I call this special case *the Bernoulli Skills Model* because each individual's skill set can be thought of as the result of a set of M Bernoulli trials, each with probability p of success. This means that the distribution of skill set sizes in the population is binomial, implying that the fraction problem solvers who have a *particular* set of k skills is $\Psi\left(A\right) = p^k \left(1-p\right)^{M-k}$, and the fraction having *any* k skills is

---

[27] To see this, consider the Shapley value of an additional skill. When an each individual requires only one collaborator to solve a problem, the individual's degree is , (the results are similar for the case where multiple collaborators may be required). Using degree, $d\left(A_i\right) = \sum_{C \subseteq A_i} \frac{\Psi(S \setminus C)}{\sum_{D \subseteq S \setminus C} \Psi(C \cup D)}$, as a value function (note that $d\left(.\right)$ satisfies both $d\left(\emptyset\right) = 0$ and superadditivity), the Shapely value of a skill, a, to individual i: $\phi_{a,i}\left(d\right) = \sum_{B \subseteq A_i \setminus \{a\}} \frac{1}{\binom{|A_i|}{|B|}} \left(\sum_{C \subseteq B} \frac{\Psi(S \setminus (C \cup a))}{\sum_{D \subseteq S \setminus (C \cup a)} \Psi((C \cup a) \cup D)}\right)$, which depends on both the existing skill population and the individual's current set of skills.

$\binom{M}{k} p^k (1-p)^{M-k}$. This special case has several characteristics which make it a useful starting point. First, because skills are uncorrelated and occur with equal frequency, individuals with the same number of skills will, in expectation, have the same degree. This means that the only variable of interest is skill set size, k. Second, because this model has only 2 parameters–M and p–it is easy to see how changes in the distribution of skills in the population affect the structure of the collaboration network.

3.2. **Degree Distribution of the Bernoulli Skills Model.**

Let Δ denote the distribution of expected degree. That is, $\Delta(d)$ is the fraction of the population who have expected degree d, where the expectation is taken over all $g \in \Gamma(\Psi)$. In this particular case, I will use a convenient shorthand: $\Delta_{M,p}$ represents the distribution of expected degree when M skills are independently distributed with probability p.

Theorem 5 states the closed form expression for the degree of an individual in the Bernoulli Skills Model.

Theorem 5. *Suppose each individual in a population faces a problem, $\omega(S)$, requiring M skills. If the skills are distributed independently with $Prob(a) = p \,\forall\, a \in S$, then the expected degree of an individual with k skills is*

$$E[d(k)] = p^M \left[\left(\frac{1-p+p^2}{p^2} \right)^k - 1 \right]$$

Proof. Since $\Sigma(\Psi) = 2^S$, every individual needs to make only one link. Thus, we can write $E[d(A)] = \sum_{C \subseteq A} \frac{\delta(C)}{\sigma(C)}$, where $\delta(C)$ is the fraction of the population who need skill set C and $\sigma(C)$ is the fraction of the population who can provide skill set C. Since the skills are independent, we can separate this sum according to the size of the skill set required. If we start with the individuals lacking exactly one skill in A_i and end with the individuals needing all of the skills, we obtain the following sum:

$$\begin{aligned}
E\left[d\left(A\right)\right] &= \sum_{i\in A}\frac{p^{M-1}\left(1-p\right)}{p}+\sum_{i,j\in A}\frac{p^{M-2}\left(1-p\right)}{p^{2}}+...+\frac{p^{M-k}\left(1-p\right)^{k}}{p^{k}} \\
&= \sum_{i=1}^{k}\binom{k}{i}\frac{p^{M-i}\left(1-p\right)^{i}}{p^{i}} \\
&= p^{M}\left[\left(\frac{1-p+p^{2}}{p^{2}}\right)^{k}-1\right]
\end{aligned}$$

□

Note that when skills are independent, an individual's degree on the network depends only on the size of her skill set, k. Therefore, in the Bernoulli Skills Model, it is appropriate to interpret the size of an individual's skill set as her "ability"–something that we cannot do in the more general case (recall Figure 2.1 in the previous section). Moreover, in the Bernoulli Skills Model, degree is a superadditive function of the number of skills a person has: $E\left[d\left(k+1\right)\right]>E\left[d\left(k\right)\right]+E\left[d(1)\right]$. This suggests a corollary to Theorem 5.

Corollary 6. *Suppose each individual in a population faces a problem, $\omega\left(S\right)$, requiring M skills. If the skills are distributed independently with $Prob\left(a\right)=p\,\forall\,a\in S$, then degree is a superadditive function of the size of her skill set, $k=\left|M\right|$. Moreover, her degree is strictly increasing in k.*

However, we still cannot price the skills individually in such a way that we capture degree, despite the fact that skills are independently distributed.

Theorem 7. *Suppose each individual in a population faces a problem, $\omega\left(S\right)$, requiring M skills . If the skills are distributed independently with $Prob\left(a\right)=p\,\forall\,a\in S$, then there exists no vector of prices, μ, such that $\sum_{a\in A}\mu_{a}=E\left[d\left(A\right)\right]$ for all $A\subseteq S$.*

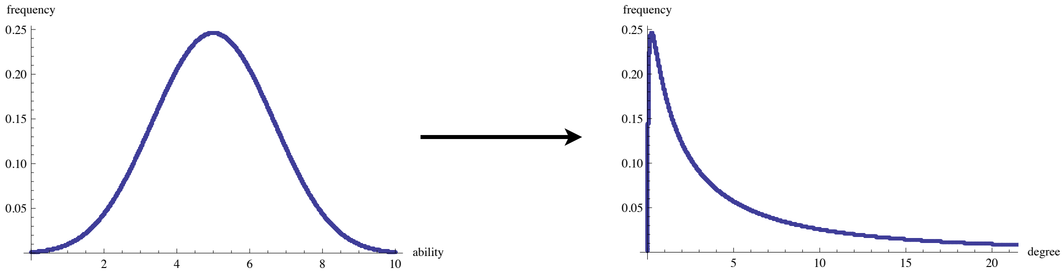

Figure 3.1. Distributions of ability and degree in a Bernoulli Skills Network with $M = 10$ and $p = \frac{1}{2}$.

Proof. Any such vector would be required to set $\mu_a = d(a) = p^{M-2}(1-p)$ for all $a \in S$. But that would imply that $d(A) = kp^{M-2}(1-p)$ for $|A| = k$. This is clearly not true for $k > 1$. $\square$

This superadditive relationship between skills and degree has a dramatic affect on the distribution of links in the population. Although the distribution of "ability" (skill set size) in the Bernoulli Skills Model is binomial, and thus symmetric, the distribution of degree is highly skewed. Individuals with a few extra skills will have many additional combinations of skills, so individuals with marginally more skills will have dramatically more links. As a result, the distribution of degree has much higher variance than the distribution of ability. Figure 3.1 shows an example of ability and degree distributions with $M = 10$ and $p = \frac{1}{2}$.

The resulting network has a very distinctive "hub and spoke" structure, in which a small number of individuals (the hubs) participate in a majority of the collaborations, while most (the spokes) participate in relatively few (see Figure 3.2 for an example network). Empirical collaboration networks tend to have a similarly skewed degree distribution. Figure 3.3 shows an example–a coauthorship network of network scientists (data courtesy of Newman (2006)).

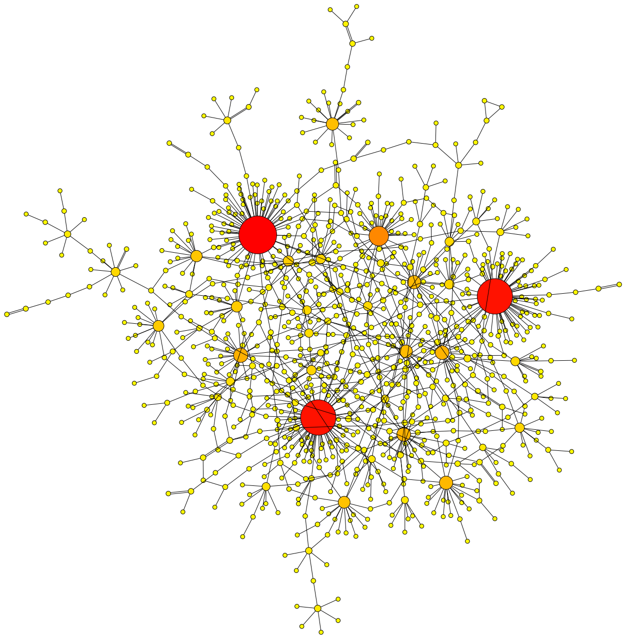

Figure 3.2. A Bernoulli Skills Network with 1000 problem solvers, $M = 16$ and $p = \frac{1}{2}$.

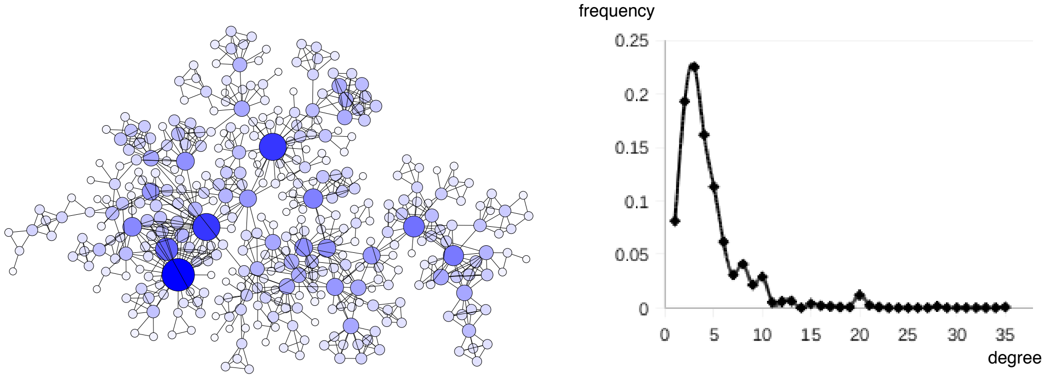

Figure 3.3. A coauthorship network of network scientists (Newman (2006)) and its degree distribution.

Simlar patterns can be observed in a wide variety of other contexts.[28] It is worth noting that this is the first decision-based model of network formation to replicate this empirical feature of social networks. Note, in particular, that while preferential attachment (a statistical model of network formation) will create a skewed degree distribution, the only distinguishing feature of high-degree nodes is their age–older nodes have a better chance of benefiting from the growth of the network.

---

[28] For example, interfirm collaboration (Powell et al. (1996), Iyer et al. (2006)), creative artists in broadway plays (Uzzi and Spiro (2005)), film actors (Barabási and Albert (1999)), jazz musicians (Gleiser and Danon (2003)), and coauthorship networks in a variety of fields: see Newman (2001) (physics), Moody (2004) (sociology), Goyal et al. (2006) (economics), and Acedo et al. (2006) (management science).

This exaggeration in the degree distribution has an empirical implication–individuals who have a disproportionate number of links in a collaboration network may not have a disproportionate number of skills. In particular, individuals who are only slightly more skillful may be dramatically more important in the collaborative community.[29] This means that collaborative communities may have "superstars"–individuals with a disproportionate influence in the community–even when skills are distributed relatively evenly in the population. Moreover, modeling individuals as having skill sets may account for some unexplained variation in collaborative outcomes.

3.3. Problem Difficulty and the Distribution of Degree: Comparative Statics on the Bernoulli Skills Model.

The hub-and-spoke structure of the collaboration network has implications for other behaviors that take place over collaborative ties–the more concentrated links are in the hands of a few hubs, the more influence they will have over opinion formation, reputation, and the decisions over new technology than they would in a network with a more equitable distribution of links. In this section, I use comparative statics of the Bernoulli skills model to examine how changes in the distribution of skills change the amount of inequality in the distribution of links.

---

[29] Thus, while great innovators like Thomas Edison are undoubtedly more skilled than the average entrepreneur, they are likely not 1000 times as skilled, as their levels of productive output might suggest.

In the following, I will use the Gini coefficient as my measure of distributional equality (See Appendix A.3 for a discussion of the gini coefficient in this case). Using the fact that $\Delta_{M,p}\left(d\left(k\right)\right) = \binom{M}{k} p^k \left(1-p\right)^{M-k}$ for $k \in \{0, 1, 2, ...M\}$, the gini coefficient of the degree distribution $\Delta_{M,p}$ is

$$G\left(p, M\right) = 1 - \frac{\left(1-p^M\right)\left(1-p\right)}{p^M\left[\left(1-p+p^2\right)\left[\left(\frac{1-p+p^2}{p^2}\right)^M - 1\right] - M\left(1-p\right)\right]}$$

As is conventional, values of $G\left(p, M\right)$ closer to 1 indicate a more unequal distribution of degree, while values closer to 0 indicate a more equitable distribution of degree.

Recall that the Bernoulli Skills model has only two parameters: the number of skills required to solve the problem (M) and the fraction of those skills possessed by the average individual (p). Figure 3.4 shows the Gini coefficient for various values of M and Figure 3.6 shows it for various values of p. It can be shown that $\frac{\delta G}{\delta M} > 0$, meaning that as problems require more skills, the distribution of degree becomes increasingly skewed towards a few high-degree nodes. Similarly, as $p \to 0$, the Gini coefficient decreases meaning that as the probability of having a given skill goes down, the network becomes more skewed (see Figure 3.6).[30]These results are summarized in Theorem 8.

Theorem 8. *Suppose each individual in a population faces a problem,* $\omega\left(S\right)$, *requiring* M *skills and further suppose the skills are distributed independently with* $Prob\left(a\right) = p\,\forall\, a \in S$. *Let* $G\left(p, M\right)$ *be the Gini Coefficient for the resulting degree distribution,* $\Delta_{M,p}$. *Then*

---

[30] Note that The behavior for p near 1 is slightly more complicated. For high values of p, many individuals in the network do not need to collaborate to solve their problem, and the network starts to break apart, resulting in an uptick in the gini coefficient for p near 1. The lower bound on this range is $p = .88$. In this paper, I am concerned with cases where collaboration is necessary, so the behavior of the gini coefficient in this upper range of p values is not particularly pertinent. More detailed information can be found in the proof of Theorem 8.

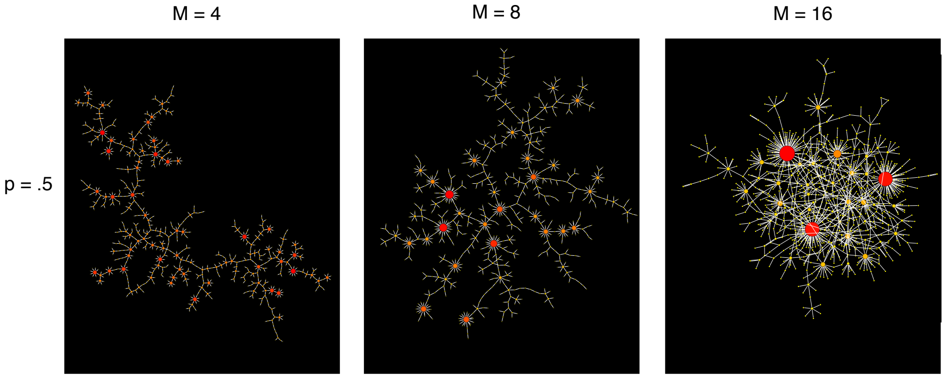

Figure 3.4. Examples of Bernoulli Skills Networks with $N = 1000$

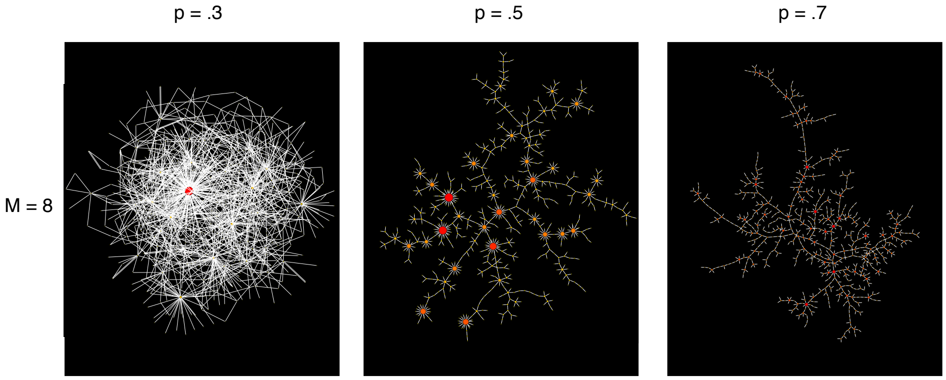

Figure 3.5. Examples of Bernoulli Skills Networks with $N = 1000$

(1) *$G(p, M)$ is strictly increasing in M. That is, the distribution of links in the collaboration network is more uneven when the problem being solved requires more skills.*

(2) *As $p \to 0$, $G(p, M)$ is strictly decreasing in p. That is, the distribution of links in the collaboration network is more uneven as the probability of having each skill gets lower.*

Proof. These relationships are illustrated in Figure 3.6. See Appendix A.4 for details. □

Together, these results can be interpreted in terms of the difficulty of the problem being faced. A problem is more difficult if it requires many skills ($M \uparrow$), or if the average individual has only a few of them ($p \downarrow$). This means that the above can be recast as a comparative static on problem difficulty. As problems become increasingly difficult ($M \uparrow$ and $p \downarrow$), the degree distribution of the network becomes more skewed, and a small number of "superstars" will dominate the collaborative community.

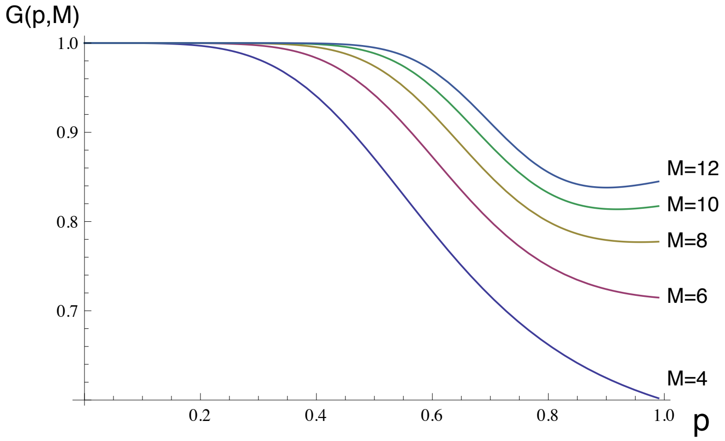

Figure 3.6. The gini coefficient, $G(p, M)$, for different values of M and p.

4. Skill Ladders: Skill Hierarchy and Degree

In the previous section, I considered a special case in which skills are entirely uncorrelated. In this section, I consider the effect of correlations between skills. I divide the skills into multiple skill ladders, where the skills within a ladder build on one another. I show that when skills are correlated in this way, the degree distribution of the collaboration network becomes even more unequal than when skills are uncorrelated. This suggests that when skills in a community are structured in this way, a very small number of individuals will tend to dominate the collaboration network.

4.1. **The Ladder Model.** In order to consider the effect of correlations between skills, in this section, I will need a second special case, which I will call the Ladder Model. I will define a *ladder* to be an ordered set of skills, $L = \{a_1, a_2, a_3...a_l\} \subseteq S$, such that $Prob\,(\text{have } a_i | \text{have } a_{i+1}) = 1$ (that is, such that any individual who has the i^{th} skill in the set must have all of the skills that precede it in the set).[31] Here, I consider a special case where the skills in S are partitioned into m ladders of equal length.[32] The set of ladders is denoted $\hat{S} = \{L_1...L_m\}$. Figure 4.1 shows an example with 12 skills arranged into four ladders.

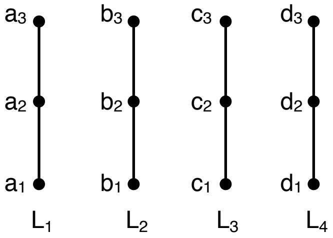

FIGURE 4.1. 12 skills arranged into four ladders of equal length.

I will call an individual who has all of the skills in a single ladder an "expert" in that ladder, and I will call the set of ladders that individual i is an expert in $\hat{A}_i \subseteq \hat{S}$.

One additional assumption will allow me to compare the results in this section to the results of the Bernoulli skills model. I will assume that the conditional probability of having the next skill in a ladder is the same for all skills–that is, $Prob\,(\text{have } a_i | \text{ have } a_{i-1}) = p$ for all a_i.[33] The probability of being an expert in a ladder of l skills is then p^l. Thanks to this assumption, the case where $m = M$ corresponds to the Bernoulli skills model. On the other extreme, $m = 1$, and all of the skills are arranged in a single ladder. Before considering ladders of arbitrary length, I will first look at this case, where $m = 1$.

---

[31] Page (2007) introduces this concept of skill ladders.

[32] Obviously, since the length of a ladder is an integer, there will only be equal-length ladders if m divides M evenly. To simplify the exposition, I have written the following as if this is true. However, all of the the following results hold if the ladders are equal length up to integer constraints, which allows for cases where m does not divide M evenly.

[33] In other words, I assume that putting a skill at the end of a ladder doesn't change the essential difficulty of obtaining that skill. If one instead makes a skill at the top easier to obtain (eg: because it builds on previous experience) or harder to obtain (eg: because they

4.2. **Example: a single ladder of skills.** Suppose the skills in the set S comprise a single ladder of length M. Because $Prob\,(\text{have } a_i | \text{have } a_{i+1}) = 1$ $\forall i \in S$, an individual's skill set can be represented by the number of skills she has ($|A_i| = k$ implies $A_i = \{a_1, a_2 ... a_k\}$) . Because individuals can be ranked in order of the number of skills they have, this case corresponds to a model in which each individual has a one-dimensional ability measure.

The linking behavior in this case is very simple. The only individuals who have skill a_M are those who also have skills $a_1 ... a_{M-1}$. Anyone who doesn't have all M skills links to someone who does. The resulting collaboration network is a set of isolated stars, each with $\frac{1-p^M}{p^M}$ links, on average. Figure 4.2 compares the network structure in the case with one skill ladder ($m = 1$) to the network structure in the Bernoulli skills model, where skills are independent ($m = M$). The populations that make up these two networks are similar–they have the same number of problem solvers, the same number of skills, and the same probability of having an additional skill. This means that the probability of an individual having all of the skills required to solve the problem is the same in both networks. Moreover, in both, exactly one individual has all of the skills required. However, the two networks have a much different structure. When all of the skills build on each other, the network is centered around a single,

---

are more demanding than the skills that came before), then the results remain qualitatively unchanged.

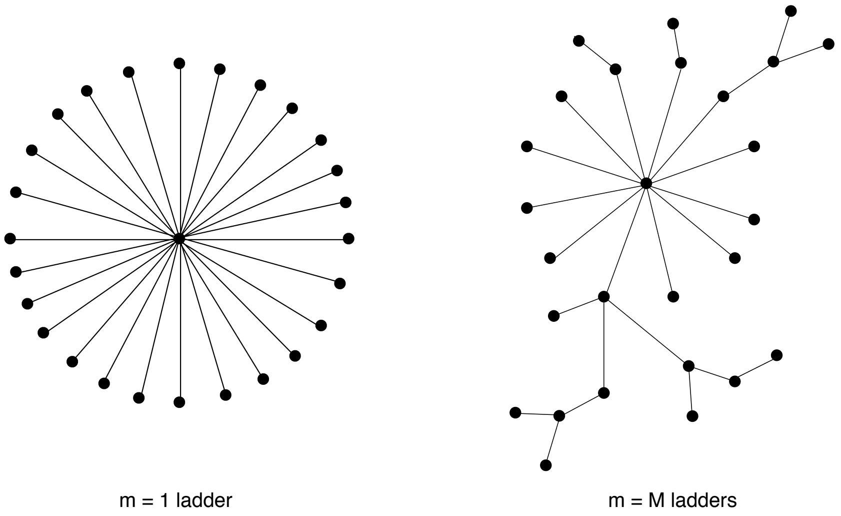

FIGURE 4.2. Ladder Skills Networks with $N = 27$ and $M = 3$.

high-degree node. When skills are independent, the network structure is much more distributed.[34]

4.3. **Results for m ladders.** Now consider the more general result. Suppose the skills are arranged in m equal-length ladders. As in the previous example, only experts obtain links. Thus, a model with m ladders reduces to a Bernoulli skills model with m independent skills. Theorem 9 presents a closed-form expression for a individual problem solver's degree in the case with m skill ladders.

---

[34] Note that this could be interpreted as a prediction about organizational structures in different industries. In industries where value is created through the exploitation of existing knowledge (eg: manufacturing), skills tend to build on one another, and it is viable to rank workers by ability. However, as demand for workers shifts towards industries that create value by generating new knowledge, we expect skills to be arranged in ladders less often. This model predicts a corresponding shift from hierarchical organizational structures towards more distributed organizational structures. Empirical evidence seems to support that prediction. Traditional organizational structures were hierarchical (the theoretical underpinnings are explored in Rosen (1982)). However, evidence indicates that organizational structure within firms is changing–hierarchical structures are flattening, and workplaces are becoming more decentralized (see, for example, Bresnahan et al. (2002) and Rajan and Wulf (2006)).

Theorem 9. *If Ψ is a distribution of skills such that $\hat{S} = \{L_1...L_m\}$ is a partition of S into m equal-length ladders with Prob $\left(\text{have } a_i^j \mid \text{ have } a_{(i-1)}^j\right) = p$, then an individual with the skill set A will have expected degree $E\left[d(A)\right] = p^M \left[\left(\frac{1-p^{\frac{M}{m}}+p^{2\frac{M}{m}}}{p^{2\frac{M}{m}}}\right)^k - 1\right]$, where k is the number of disciplines the individual is an expert in.*

Proof. The ladders are of equal length, so the length of a single ladder is $\frac{M}{m}$, and the probability that an individual is an expert in any one ladder is $p^{\frac{M}{m}}$. An individual receives a link only if she is an expert in a field. Define a new set of skills that correspond to the set of ladders: $\hat{S} = \{L_1...L_m\}$. The individual's new skill set is $\hat{A}_i$, where $L_k \in \hat{A}_i$ if she is an expert in ladder L_k. Each of these new skills has a probability equal to the probability of being an expert in that field, so define $\hat{p} = p^{M/m}$. The probability of being an expert in a particular ladder is independent of the probability of being an expert in any other ladder, so this problem reduces to one with m independent skills, with probability $p^{\frac{M}{m}}$. The result then is a simple extension of Theorem 5. $\square$

We can now do a comparative static on the number of skill ladders, to see how the amount of hierachy in the skill pool affects the structure of the collaboration network. Theorem 10 indicates that in communities where skills are more hierachical, the degree distribution of the collaboration network is more skewed, and a few experts become extremely high-degree hubs in the network. Figure 4.3 shows how the gini coefficient depends on the number of ladders for the case where $M = 6$ (note that the bottom curve ($m = 6$) is equivalent to the Bernoulli Skills Model with $M = 6$ in Figure 3.6).

Theorem 10. *Suppose S skills are arranged in m ladders of equal length, with constant conditional probability Prob $\left(\text{have } a_i^j \mid \text{ have } a_{i-1}^j\right) = p$ and Prob $\left(\text{have } a_1^j\right) = p \; \forall j = 1...m$. The gini coefficient of the resulting network is decreasing in the*

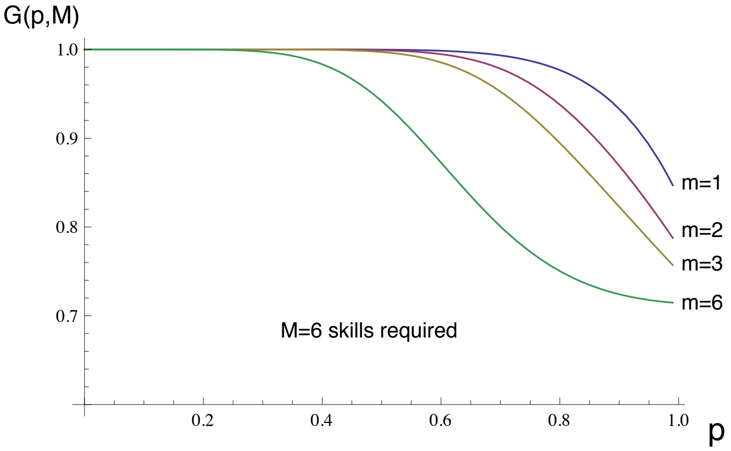

Figure 4.3. The gini coefficient, $G\left(p, M\right)$, for different values of m and p.

number of ladders, m. That is, when there are fewer skill ladders, the degree distribution becomes increasingly uneven.

Proof. A model with m skill ladders is equivalent to a Bernoulli Skills model with m skills and $p = p^{M/m}$. The gini coefficient is thus $G\left(p, \frac{M}{m}\right)$. Since this function is increasing in the second argument (see Theorem 8), it is decreasing in m. $\square$

5. Discussion: The Importance of Heterogeneity

One of the main contributions of the model presented above is it's extremely general treatment of individual heterogeneity. Although it is possible to model the heterogeneity of problem solvers more simply–for example, via a one-dimensional ability level,[35] by giving each individual a type or speciality,[36] or some combination of the two[37]–it is not clear that those models capture what distinguishes one problem solver from another. The model presented in this paper encompasses both ability-based and type-based models,[38] while also allowing for other, more complex types of heterogeneity. This finer-grained approach has significant impacts on outcomes. In this section, I examine those impacts by directly comparing a model in which individuals have "types" to one in which individuals have overlapping sets of skills.

---

[35] Gautier (2002), Shi (2002)
[36] Hamilton et al. (2000)
[37] Jovanovic (1994)

Suppose there are two different problem solving populations, faced with the same problem, $\omega = \{a, b, c\}$. In population A, every individual has exactly one skill, which we can think of as her speciality or "type": $\Psi(a) = \frac{1}{3}$, $\Psi(b) = \frac{1}{3}, \Psi(c) = \frac{1}{3}$. In population B, the skills are distributed independently. This means that and the probability of having skill set A is $\Psi(A) = \prod_{i \in A} p_i$. Let $p_i = Prob\,(have\,skill\,i) = \frac{1}{3}$ for $i = a, b, c$.

These two populations have much in common: in both, the three skills occur in equal proportion, and both populations average one skill per person. However, as illustrated in Figure 5.1, the collaboration network is much different in the two different populations. This pair of examples highlights the value of a more detailed treatment of skills in modeling problem solving. The assumption that every individual has a type or specialty is not necessarily benign–it impacts our predictions about the value of certain individuals in the community, as well as impacting the structure of the collaboration network.

6. Discussion of Implications, Extensions, and Conclusion

In this paper, I have presented a model of collaboration network formation in which individuals have heterogeneous skill sets and collaborate to solve problems that none of them could solve as individuals. The result is a collaboration network, the structure of which depends on the distribution of skills in the underlying community. I have showed that when individuals have complex

---

[38] As mentioned earlier, a ladder model with $m = 1$ ladder is equivalent to an ability measure. A type-based model is equivalent to a population where each individual a single skill (as in the below example).

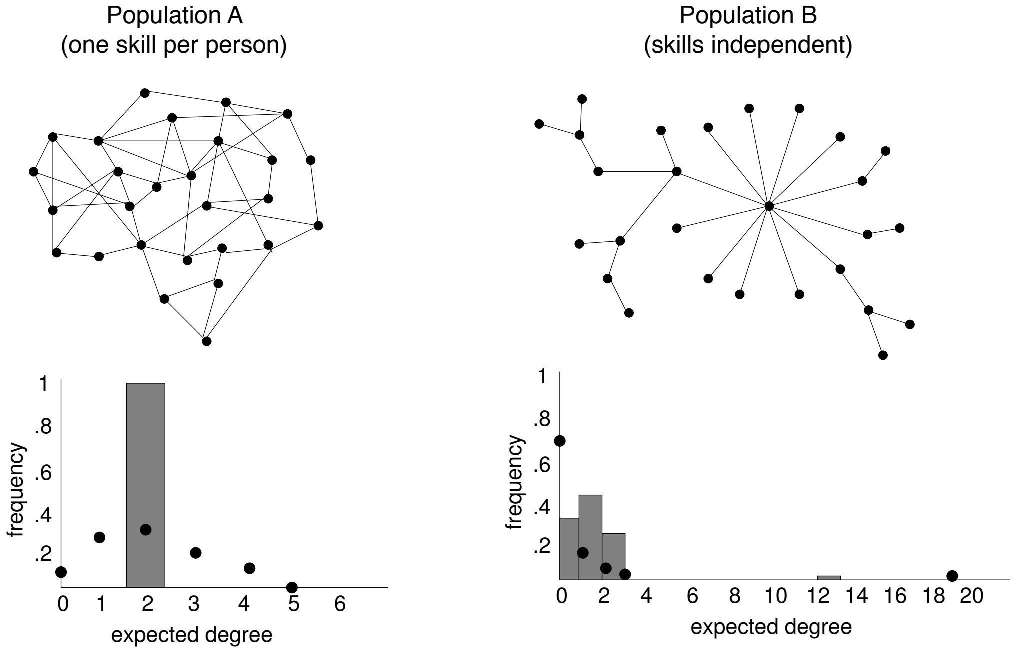

Figure 5.1. A population with types and a population with independent skills.

skill sets, their combination of skills is important, meaning that their degree in the network cannot be captured via a linear pricing system. Using a special case of the model in which skills are distributed independently, I show that small initial differences in skill sets are magnified in the network into large differences in degree. This means that the network can have a skewed degree distribution, even when skills are distributed symmetrically. This corresponds with an empirical fact: the degree distribution of collaboration networks is highly skewed, despite relatively low variation in the perceived distribution of ability. Finally, I show that as problems become more difficult, the distribution of degree in the network becomes increasingly skewed towards a few, high-degree "super stars".

The more general treatment of individual heterogeneity that I present in this paper has some interesting implications for labor markets knowledge-based fields, where workers are richly heterogeneous and collaborate to solve problems.[39] In particular, given that there is a relationship between the number of collaborators an individual has and her output,[40] the model presented in this paper can provide insight into individual heterogeneity and output. The results of Section 5 strongly indicate that the decisions we make when modeling heterogeneity matter. A full exploration of the effects of complex heterogeneity on labor markets is outside the scope of this work, but I will provide a brief discussion here.

Empirically, output in creative and knowledge-based fields is highly concentrated among a small number of people.[41] This long-tailed distribution of productive output has implications for the distribution of wages and welfare, and thus there has been considerable interest in understanding why such a concentration in labor demand occurs. Rosen (1981)'s model of production can induce a long-tailed distribution when there is a high premium on quality, and production technology decouples effort from output quantity (eg: in creative industries, where a single performance or album can be enjoyed by many consumers). However, such technologies are not relevant in knowledge-based industries, where effort is not decoupled from output volume. The long-tailed distribution of collaborative interactions in this paper suggests an alternative explanation for the observed variation in output. Moreover, the prediction that output will become more skewed as problems become more difficult suggests an empirical test in the context of aggregate labor demand.

---

[39] If we were to recast the relationship between collaborators as a relationship between employees and firms, then an extension of this model would also provide insights into the affect of heterogeneity on labor in a non-collaborative context.

[40] In the Bernoulli Skills model, where each individual has a complement, the relationship is one-to-one.

[41] Newman (2001), Moody (2004), Acedo et al. (2006), and Goyal et al. (2006) show that a small number of academics produce the majority of papers written, Rosen (1981) observes a similar pattern in in music, film, and textbook writing, and Uzzi and Spiro (2005) notes a similarly skewed distribution of output among directors, producers, and other creative artists on Broadway.

The non-linear relationship between skill sets and output also has implications for the explanation of variation in labor outcomes when individual heterogeneity is imperfectly observed. If we assume that skills make a linear contribution to output, then we would expect a linear decline in the amount of variation explained. However, if the relationship between individual skills and output is non-linear, as is suggested by the model in this paper, then the amount of variation in output explained by observables will drop nonlinearly as well. This suggests that decreases in the amount of variation in labor outcomes explained by observed agent heterogeneity could be the result of a shift from manufacturing–where skill could reasonably be considered to be a one-dimensional variable–to knowledge-based industries, where individual heterogeneity is more complex. A more thorough exploration of the implications of skill-based labor heterogeneity on labor markets would make an interesting extension.

Turning to future work, the framework that I have introduced in this paper is extremely flexible and has the potential to provide insights into a wide variety of questions that could not previously be answered. Moreover, the results of this paper are evidence that more complex models of individual heterogeneity are both tractable and can provide the means for addressing a wide range of new and important questions surrounding problem solving and innovation.

References

Acedo, Francisco Jose, Carmen Barroso, Cristobal Casanueva, and Jose Luis Galan, "Co-authorship in management and organizational studies: an empirical and network analysis," *Journal of Management Studies*, 2006, *43* (5), 957–983.

Astebro, Thomas and Peter Thompson, "Entrepreneurs, Jacks of all trades or Hobos?," *Research Policy*, 2011, *40* (5), 637–649.

Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson, "The Diffusion of Microfinance," 2011.

Barabasi, A-L, H Jeong, Z Neda, E Ravasz, A Schubert, and T Vicsek, "Evolution of the social network of scientific collaborations," *Physica A: Statistical Mechanics and its Applications*, August 2002, *311* (3-4), 590–614.

Barabási, Albert-László and Réka Albert, "Emergence of Scaling in Random Networks," *Science*, October 1999, *286* (5439), 509–512.

Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Calvó-Armengol, Antoni, "Job contact networks," *Journal of Economic Theory*, March 2004, *115* (1), 191–206.

_______ **and Matthew O Jackson**, "The Effects of Social Networks on Employment and Inequality," *American Economic Review*, 2004, *94* (3), 426–454.

DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion bias, social influence, and uni-dimensional opinions," *Quarterly Journal of Economics*, 2003, *118* (3), 909–968.

Erdös, Paul and Alfréd Rényi, "On random graphs," *Publ Math Debrecen*, 1959, *6*, 290–297.

Gautier, Pieter A, "Unemployment and Search Externalities in a Model with Heterogeneous Jobs and Workers," *Economica*, 2002, *69* (273), 21–40.

Gleiser, Pablo and Leon Danon, "Community Structure in Jazz," *Advances in Complex Systems*, 2003, *6* (4), 565–573.

Golub, Benjamin and Matthew O Jackson, "Naïve Learning in Social

Networks and the Wisdom of Crowds," *American Economic Journal Microeconomics*, 2010, *2* (1), 112–149.

Goyal, Sanjeev and José Luis Moraga-González, "R&D Networks," *The RAND Journal of Economics*, 2001, *32* (4), 686–707.

_______**, Marco J Van Der Leij, and Jose Luis Moraga-Gonzalez**, "Economics: An Emerging Small World," *Journal of Political Economy*, 2006, *114* (2), 403–412.

Grossman, Jerrold W and Patrick D F Ion, "On a Portion of the Well-Known Collaboration Graph," *Congressus Numerantium*, 1995, *108*, 129–131.

Guimerà, Roger, Brian Uzzi, Jarrett Spiro, and Luís A Nunes Amaral, "Team assembly mechanisms determine collaboration network structure and team performance.," *Science*, 2005, *308* (5722), 697–702.

Hamilton, Jonathan, Jacques-François Thisse, and Yves Zenou, "Wage competition with heterogeneous workers and firms," *Journal of Labor Economics*, 2000, *18* (3), 453–472.

Hong, Lu and Scott E Page, "Problem Solving by Heterogeneous Agents," *Journal of Economic Theory*, 2001, *97* (1), 123–163.

Iyer, Bala, Chi-Hyon Lee, and N Venkatraman, "Managing in a Small World Ecosystem : Some Lessons from the Software Sector," *California Management Review*, 2006, *48* (3), 28–47.

Jackson, Matthew O., *Social and Economic Networks*, Princeton University Press, 2008.

Jackson, Matthew O and Asher Wolinsky, "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 1996, *71* (1), 44–74.

_______ **and Brian W Rogers**, "Relating Network Structure to Diffusion Properties through Stochastic Dominance," *The BE Journal of Theoretical Economics*, 2007, *7* (1), 1–13.

——— **and Leeat Yariv**, "Diffusion of Behavior and Equilibrium Properties in Network Games," *American Economic Review*, 2007, *97* (2), 92–98.

Jovanovic, Boyan, "Firm formation with heterogeneous management and labor skills," *Small Business Economics*, 1994, *6* (3), 185–191.

Laband, David N and Robert D Tollison, "Intellectual Collaboration," *Journal of Political Economy*, 2000, *108* (3), 632–662.

Lazear, Edward P, "Balanced Skills and Entrepreneurship," *American Economic Review*, 2004, *94* (2), 208–211.

———, "Entrepreneurship," *Journal of Labor Economics*, 2005, *23* (4), 649–680.

Menzel, Herbert and Elihu Katz, "Social Relations and Innovation in the Medical Profession: The Epidemiology of a New Drug," *The Public Opinion Quarterly*, 1955, *19* (4), 337–352.

Moody, James, "The Structure of a Social Science Collaboration Network: Disciplinary Cohesion from 1963 to 1999," *American Sociological Review*, 2004, *69* (2), 213–238.

Newman, M.E.J., "Scientific collaboration networks. I. Network construction and fundamental results," *Physical Review E*, 2001, *64* (1), 1–8.

———, "No Title," *Physical Review E*, 2006, *74*, 036104.

Page, Scott E., *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press, 2007.

Phillips, Katherine W, Elizabeth A Mannix, Margaret A Neale, and Deborah H Gruenfeld, "Diverse groups and information sharing: The effects of congruent ties," *Journal of Experimental Social Psychology*, 2004, *40* (4), 497–510.

Polzer, Jeffrey T, Laurie P Milton, and William B Swann, "Capitalizing on diversity: Interpersonal congruence in small work groups," *Administrative Science Quarterly*, 2002, *47* (2), 296–324.

Powell, Walter W, Kenneth W Koput, and Laurel Smith-Doerr, "Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology," *Administrative Science Quarterly*, 1996, *41* (1), 116.

Rajan, Raghuram G and Julie Wulf, "The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies," *The Review of Economics and Statistics*, 2006, *88* (4), 759–773.

Ramasco, Jose J, S N Dorogovtsev, and Romualdo Pastor-Satorras, "Self-organization of collaboration networks," *Physical Review E*, 2007, *70* (3), 10.

Rosen, Sherwin, "The Economics of Superstars," *American Economic Review*, 1981, *71* (5), 845–858.

———, "Authority, Control, and the Distribution of Earnings," *The Bell Journal of Economics*, 1982, *13* (2), 311–323.

Roy, A D, "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 1951, *3* (2), 135–146.

Shi, Shouyong, "A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Biased Technology," *The Review of Economic Studies*, 2002, *69* (2), 467–491.

Thomas-Hunt, Melissa C, Tonya Y Ogden, and Margaret A Neale, "Who's Really Sharing? Effects of Social and Expert Status on Knowledge Exchange Within Groups," *Management Science*, 2003, *49* (4), 464–477.

Uzzi, Brian and Jarrett Spiro, "Collaboration and Creativity: The Small World Problem," *American Journal of Sociology*, 2005, *111* (2), 447–504.

Watts, Duncan J and Steven H Strogatz, "Collective dynamics of 'small-world' networks.," *Nature*, 1998, *393* (6684), 440–2.

APPENDIX A.

A.1. **Pairwise Stability and Efficiency.**

Briefly, a network is pairwise stable if no individual would prefer to terminate an existing link, and if no pair of individuals would prefer to add a link. Although this definition is usually used in undirected networks, it works equally well in the current context. Formally, a directed collaboration network, g, is *pairwise stable* if

1. for all $ij \in g$, $u_i(g) \geq u_i(g - ij)$ and $u_j(g) \geq u_j(g - ij)$
2. for all $ij \notin g$, if $u_j(g + ij) > u_j(g)$ then $u_i(g + ij) < u_i(g)$

Together, these two conditions ensure that links are mutual. That is, if a network is pairwise stable, then both ends agree to maintain the link.

Theorem. *Any complementarity network, $g \in \Gamma(\Psi)$, is pairwise stable. In other words, $\forall ij \in g$ $u_i(g) \geq u_i(g - ij)$ and $u_j(g) \geq u_j(g - ij)$ and for all $ij \notin g$, if $u_j(g + ij) > u_j(g)$ then $u_i(g + ij) < u_i(g)$.*

Proof. First, consider whether any individual wishes to unilaterally remove a link, $ij \in g$. Removing this incoming link costs individual j her share of the payoff from solving $i's$ problem $\left(\frac{1}{|C_i|+1} \geq 0\right)$, and thus she will never choose to do so. Individual i chose a minimal set of collaborators that allowed her to solve her problem, so removing this outgoing link means that she can no longer solve the problem. This would cost her the payoff $\left(\frac{1}{|C_i|+1} \geq 0\right)$ from solving the problem, and thus she will also never choose to do so. Finally, note that no individual would ever want to add an outgoing link to a cost-minimizing collaboration network, because having chosen a minimal set of collaborators, any additional link would require her to further split her prize. □

Theorem. *Any cost minimizing collaboration network, $g \in \Gamma(\Psi)$, is strongly efficient. In other words, $\sum_i u_i(g) \geq \sum_i u_i(g') \; \forall g' \in G$.*

Proof. Because all value is generated from solving problems, the maximum possible value in the network is N. Since solving problems is incentive compatible and there is no loss, the problem solvers always extract the maximum value from the network. $\square$

A.2. **General Proof of Theorem 2.** An individual with the set $A\cup B$ will be able to help anyone needing any subset of those skills. Let $\delta\left(C\right)$ be the demand for a particular set of skills, C. In the general case, $\delta\left(C\right)=\Psi\left(S\backslash C\right)+\sum_{D\,:\,\Psi(C\cup D)=0}\Psi\left(S\backslash\left(C\cup D\right)\right)$. The fraction who can supply the set C is $\sigma\left(C\right)=\sum_{D\subseteq S\backslash C}\Psi\left(C\cup D\right)$. Note that $\delta\left(C\right)$ and $\sigma\left(C\right)$ depend only on the particulars of the problem (S), the distribution of skills (Ψ), and the subset of skills (C). Thus, any individual with the skill set $A\cup B$ has expected degree

$$E\left[d\left(A\cup B\right)\right]=\sum_{C\subseteq A\cup B}\frac{\delta\left(C\right)}{\sigma\left(C\right)}$$

We can divide the problems that an individual with $A\cup B$ can solve into three groups:

(1) Requires only skills from set A: $C\subseteq A$
(2) Requires only skills from set B, including at least one found only in B : $\{C\,|\,C\subseteq B\,and\,\exists\,b\in C\,st\,b\in B\backslash A\}$
(3) Requires at least one skill from each set that can only be found in that set: $\{C\,|\,C\subseteq A\cup B,\,where\,\exists\,a,b\in C\,s.t.\,a\in A\backslash B\,and\,b\in B\backslash A\}$

Using this partition, we can write

$$\begin{aligned}E\left[d\left(A\cup B\right)\right] &= \sum_{C\subseteq A}\frac{\delta\left(C\right)}{\sigma\left(C\right)}+\sum_{C\subseteq B\,and\,C\cap B\neq\emptyset}\frac{\delta\left(C\right)}{\sigma\left(C\right)}+\sum_{C\subseteq A\cup B\,and\,C\cap A,\,C\cap B\neq\emptyset}\frac{\delta\left(C\right)}{\sigma\left(C\right)}\\ &= E\left[d\left(A\right)\right]+\sum_{C\subseteq B\,and\,C\cap B\neq\emptyset}\frac{\delta\left(C\right)}{\sigma\left(C\right)}+\phi\end{aligned}$$

which implies that

$$\begin{aligned}
E\left[d\left(A\cup B\right)\right]+E\left[d\left(A\cap B\right)\right] &= E\left[d\left(A\right)\right]+\sum_{C\subseteq B\, and\, C\cap B\neq\emptyset}\frac{\delta\left(C\right)}{\sigma\left(C\right)}+\phi+E\left[d\left(A\cap B\right)\right)] \\
&= E\left[d\left(A\right)\right]+\left(\sum_{C\subseteq B\, and\, C\cap B\neq\emptyset}\frac{\delta\left(C\right)}{\sigma\left(C\right)}+\sum_{C\subseteq A\cap B}\frac{\delta\left(C\right)}{\sigma\left(C\right)}\right)+\phi \\
&= E\left[d\left(A\right)\right]+E\left[d\left(B\right)\right]+\phi \\
&\geq E\left[d\left(A\right)\right]+E\left[d\left(B\right)\right]
\end{aligned}$$

A.3. **Discussion of the Gini Coefficient in the discrete case.** The Gini coefficient measures the area between the Lorenz curve of a distribution (in this case, the distribution of expected degree), and the line of equality. In the case of a discrete distribution with values $y_0 ... y_N$ where $y_i < y_{i+1}$, the Lorenz curve is a piecewise function connecting points (F_i, D_i) where $F_i = \sum_{k=0}^{i}\Delta\left(y_k\right)$ is the fraction of individuals with strictly less than y_i links, and $D_i = \frac{\sum_{k=0}^{i}\Delta(y_k)y_k}{\sum_{k=0}^{N}\Delta(y_k)y_k}$ is the fraction of the total number of links they hold. See Figure A.1 for an example. The gini coefficient for a discrete distribution is given by $G = 1-\sum_{i=1}^{N}D_i\left(F_i-F_{i-1}\right)$. Lower values of the gini coefficient indicate a more equal distribution of links across individals in the population, and higher values indicate a more skewed distribution of links. The coefficient is which is 0 when the distribution is perfectly equal (ie: the bottom $x\%$ of the population holds exactly $x\%$ of the links) and 1 when all of the links are held by a single individual.

A.4. **Proof of Theorem 8.** Recall that since $\Delta_{M,p}\left(d\left(k\right)\right)=\binom{M}{k}p^k\left(1-p\right)^{M-k}$ for $k\in\{0,1,2,...M\}$, the gini coefficient of the degree distribution $\Delta_{M,p}$ is

$$G\left(p,M\right)=1-\frac{\left(1-p^M\right)\left(1-p\right)}{p^M\left[\left(1-p+p^2\right)\left[\left(\frac{1-p+p^2}{p^2}\right)^M-1\right]-M\left(1-p\right)\right]}$$

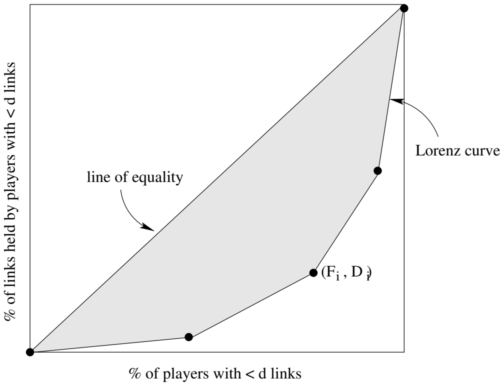
% of links held by players with < d links
line of equality
Lorenz curve
(F_i , D_i)
% of players with < d links

Figure A.1. An example of the Gini coefficient for a discrete distribution, $\Delta(y)$. In this case, the random variable y takes on one of five values, $y_0 ... y_4$. The Gini coefficient is the area of the shaded region between the line of equality and the Lorenz curve.

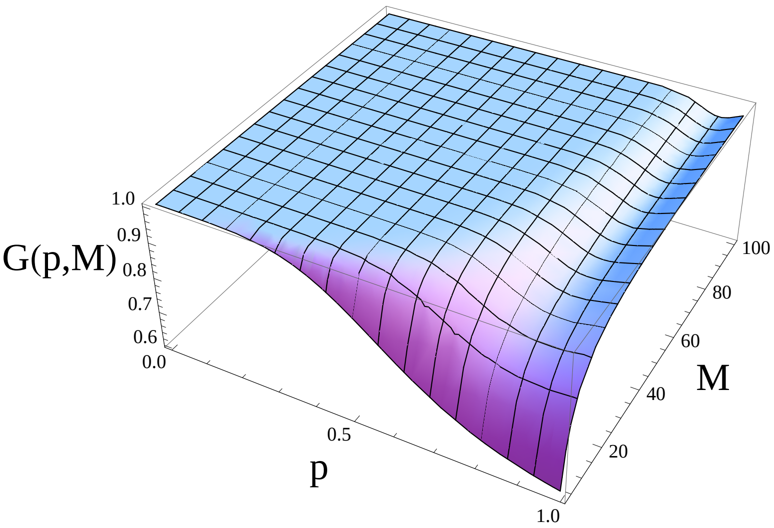

Figure A.2. Plot of the Gini Coefficient for values of M and p.

Figure A.2 shows the Gini Coefficent of $\Delta_{M,p}$ for values of M and p.

First, show that $G(p, M+1) - G(p, M) > 0$.

$$G(p, M+1) - G(p, M) = \frac{(1-p)^2\Big[\Big(\big(1-p+p^2\big)^{M+1}-p^{2M+2}\Big)\big(1-p-p^M\big)+p^{2M+1}(M+1)(1-p)\Big]}{p^{3M}\Big[-(M+1)(1-p)-p^2+\Big(\frac{1-p+p^2}{p^2}\Big)^{M}(1-p+p^2)\Big]\Big[\Big(\frac{1-p+p^2}{p^2}\Big)^{M}(1-p+p^2)^2-p^2((2-M)(1-p)+p^2}$$

Thus, $G(p, M+1) - G(p, M) > 0$ if $\left(\left(1-p+p^2\right)^{M+1} - p^{2M+2}\right)\left(1-p-p^M\right) + p^{2M+1}(M+1)(1-p) > 0$. Since $(1-p) > \left(1-p-p^M\right)$, this is true if

$\left((1-p+p^2)^{M+1} - p^{2M+2}\right)\left(1-p-p^M\right) + p^{2M+1}(M+1)\left(1-p-p^M\right) > 0$, which is true if $(1-p+p^2)^{M+1} - p^{2M+2} + p^{2M+1}(M+1) > 0$. Since $(1-p+p^2) > p$ for $p \in [0,1]$, the previous is true if $p^{M+1} - p^{2M+2} + p^{2M+1}(M+1) > 0$, which is true if $1 - p^{M+1} + p^M(M+1) > 0$. Since $1 - p^{M+1} + p^M(M+1) = 1 + p^M(1-p) + Mp^M > 0$, the result holds.

Now, show that $G(p,M)$ is strictly decreasing in p as $p \to 0$.

Note that $G(p,M)$ is strictly decreasing in p for all $p < T(p,M)$ where $\mathrm{T}(p,M)$ is a boundary condition given by

$$M\left(-1+2p-2p^2+p^3+\left(\frac{1-p+p^2}{p^2}\right)^M\left((2-p)(1-p)p^M-(1-p)^2(1+p)\right)-M(1-p)^2\right) = \left(1-\left(\frac{1-p+p^2}{p^2}\right)^M\right)(2-p)\left(1-p^M\right)p^2$$

Through numerical analysis, $T(p,M) < .88$ for all (p,M). Thus, $G(p,M)$ is decreasing in p for $p < .88$, and the result follows